CLAIMS

What is claimed is:

1. A method for generating a speech recognition database comprising:

generating a latent semantic analysis (LSA) space from a training corpus of documents representative of a language;

receiving a new document that represents a change in the language; and

adapting the LSA space to reflect the change in the language.

2. The method of claim 1, wherein adapting the LSA space to reflect the change in the language comprises transforming the LSA space to take into account the new document's influence on the LSA space without re-computing the LSA space.

3. The method of claim 1, wherein transforming the LSA space comprises:

obtaining a training document vector that characterizes a semantic position of the training document within the LSA space;

computing a new document vector that characterizes a semantic position of the new document within the LSA space;

deriving a document vector transformation matrix; and

applying the document vector transformation matrix to the training document vector and the new document vector to shift a position of each document vector in the LSA space, where the shift in the position reflects the change in the language.

4. The method of claim 3, further comprising:

obtaining a training word vector that characterizes a semantic position of the training word within the LSA space;

computing a new word vector that characterizes a semantic position of the new word within the LSA space;

deriving a word vector transformation matrix; and

applying the word vector transformation matrix to the training word vector and the new word vector to shift a position of each word vector in the LSA space, where the shift in the position reflects the change in the language.

5. The method of claim 4, wherein:

the training document vector is VS, where VS is computed from a right singular matrix V and a diagonal matrix S, each of which was obtained from a previous singular value decomposition (SVD) of a training word-document matrix constructed during the generation of the LSA space, the training word-document matrix representing the extent to which each of the words appears in each of the documents of the training corpus;

the new document vector ZS, where ZS is computed from the diagonal matrix S and an extension matrix Z, wherein Z is an extension of the right singular matrix V obtained by folding in a new word-document matrix, the new word-document matrix representing the extent to which a new word appears in the new document; and

the document vector transformation matrix is J, wherein J is obtained from a Choleski decomposition of a matrix derived from an extension matrix Y, wherein Y is an

extension of a left singular matrix U obtained by folding in the new word-document matrix, and wherein U was obtained from the previous SVD of the training word-document matrix constructed during the generation of the LSA space.

6. The method of claim 5, wherein:

the training word vector is US, wherein US is computed from the left singular matrix U and the diagonal matrix S;

the new word vector is YS, wherein YS is computed from the diagonal matrix S and the extension matrix Y; and

the word vector transformation matrix is K, wherein K is obtained from a Choleski decomposition of a matrix derived from the extension matrix Z.

7. The method of claim 6, wherein transforming the LSA space comprises applying the document vector transformation matrix and the word vector transformation matrix simultaneously.

8. The method of claim 6, wherein when the new document matrix contains more new documents than new words, then transforming the LSA space comprises:

applying the word vector transformation matrix K, first; and

applying the document vector transformation matrix J second, wherein the extension matrix Y is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Z.

9. The method of claim 6, wherein when the new document matrix contains more new words than new documents, then transforming the LSA space comprises:

applying the document vector transformation matrix J first; and

applying the word vector transformation matrix K second, wherein the extension matrix Z is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Y.

10. The method of claim 1, wherein the change in the language is a change in the language's domain.

11. The method of claim 1, wherein the change in the language is a change in the language's style.

12. A computer-readable medium having executable instructions to cause a computer to perform a method for generating a speech recognition database comprising:

generating a latent semantic analysis (LSA) space from a training corpus of documents representative of a language;

receiving a new document that represents a change in the language; and

adapting the LSA space to reflect the change in the language.

13. The computer-readable medium of claim 12, wherein adapting the LSA space to reflect the change in the language further comprises transforming the LSA space to take

into account the new document's influence on the LSA space without re-computing the LSA space.

14. The computer-readable medium of claim 13, wherein transforming the LSA space further comprises:

obtaining a training document vector that characterizes a semantic position of the training document within the LSA space;

computing a new document vector that characterizes a semantic position of the new document within the LSA space;

deriving a document vector transformation matrix; and

applying the document vector transformation matrix to the training document vector and the new document vector to shift a position of each document vector in the LSA space, where the shift in the position reflects the change in the language.

15. The computer-readable medium of claim 14, wherein transforming the LSA space further comprises:

obtaining a training word vector that characterizes a semantic position of the training word within the LSA space;

computing a new word vector that characterizes a semantic position of the new word within the LSA space;

deriving a word vector transformation matrix; and

applying the word vector transformation matrix to the training word vector and the new word vector to shift a position of each word vector in the LSA space, where the shift in the position reflects the change in the language.

16. The computer-readable medium of claim 14, wherein:

the training document vector is VS where VS is computed from a right singular matrix V and a diagonal matrix S, each of which was obtained from a previous singular value decomposition (SVD) of a training word-document matrix constructed during the generation of the LSA space, the training word-document matrix representing the extent to which each of the words appears in each of the documents of the training corpus;

the new document vector is ZS where ZS is computed from the diagonal matrix S and an extension matrix Z, wherein Z is an extension of the right singular matrix V obtained by folding in a new word-document matrix, the new word-document matrix representing the extent to which a new word appears in the new document; and

the document vector transformation matrix is J, wherein J is obtained from a Choleski decomposition of a matrix derived from an extension matrix Y, wherein Y is an extension of a left singular matrix U obtained by folding in the new word-document matrix, and wherein U was obtained from the previous SVD of the training word-document matrix constructed during the generation of the LSA space.

17. The computer-readable medium of claim 16, wherein:

the training word vector is US, wherein US is computed from the left singular matrix U and the diagonal matrix S;

the new word vector is YS, wherein YS is computed from the diagonal matrix S and the extension matrix Y; and

the word vector transformation matrix is K, wherein K is obtained from a Choleski decomposition of a matrix derived from the extension matrix Z.

18. The computer-readable medium of claim 17, wherein transforming the LSA space further comprises applying the document vector transformation matrix and the word vector transformation matrix simultaneously.

19. The computer-readable medium of claim 17, wherein, when the new document matrix contains more new documents than new words, transforming the LSA space further comprises:

applying the word vector transformation matrix K, first; and

applying the document vector transformation matrix is J second, wherein the extension matrix Y is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Z.

20. The computer-readable medium of claim 17, wherein, when the new document matrix contains more new words than new documents, transforming the LSA space comprises:

applying the document vector transformation matrix J first; and

applying the word vector transformation matrix K second, wherein the extension matrix Z is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Y.

21. The computer-readable medium of claim 12, wherein the change in the language is a change in the language's domain.

22. The computer-readable medium of claim 12, wherein the change in the language is a change in the language's style.

23. An apparatus for generating a speech recognition database, the apparatus comprising:

a latent semantic analysis (LSA) space generator to generate an LSA space from a training corpus of documents representative of a language;

a document receiver to receive a new document that represents a change in the language; and

an LSA space adapter to adapt the LSA space to reflect the change in the language.

24. The apparatus of claim 23, wherein LSA space adapter transforms the LSA space to take into account the new document's influence on the LSA space without re-computing the LSA space.

25. The apparatus of claim 23, wherein the LSA space adapter transforms the LSA space by:

obtaining a training document vector that characterizes a semantic position of the training document within the LSA space;

computing a new document vector that characterizes a semantic position of the new document within the LSA space;

deriving a document vector transformation matrix; and

applying the document vector transformation matrix to the training document vector and the new document vector to shift a position of each document vector in the LSA space, where the shift in the position reflects the change in the language.

26. The apparatus of claim 25, wherein the LSA space adapter further transforms the LSA space by:

obtaining a training word vector that characterizes a semantic position of the training word within the LSA space;

computing a new word vector that characterizes a semantic position of the new word within the LSA space;

deriving a word vector transformation matrix; and

applying the word vector transformation matrix to the training word vector and the new word vector to shift a position of each word vector in the LSA space, where the shift in the position reflects the change in the language.

27. The apparatus of claim 26, wherein:

the training document vector is VS, where VS is computed from a right singular matrix V and a diagonal matrix S, each of which was obtained from a previous singular value decomposition (SVD) of a training word-document matrix constructed during the generation of the LSA space, the training word-document matrix representing the extent to which each of the words appears in each of the documents of the training corpus;

the new document vector ZS, where ZS is computed from the diagonal matrix S and an extension matrix Z, wherein Z is an extension of the right singular matrix V obtained by folding in a new word-document matrix, the new word-document matrix representing the extent to which a new word appears in the new document; and

the document vector transformation matrix is J, wherein J is obtained from a Choleski decomposition of a matrix derived from an extension matrix Y, wherein Y is an extension of a left singular matrix U obtained by folding in the new word-document matrix, and wherein U was obtained from the previous SVD of the training word-document matrix constructed during the generation of the LSA space.

28. The apparatus of claim 26, wherein:

the training word vector is US, where US is computed from a left singular matrix U and the diagonal matrix S;

the new word vector is YS, where YS is computed from the diagonal matrix S and the extension matrix Y; and

the word vector transformation matrix is K, wherein K is obtained from a Choleski decomposition of a matrix derived from the extension matrix Z.

29. The apparatus of claim 26, wherein the LSA space adapter transforms the LSA space by applying the document vector transformation matrix and the word vector transformation matrix simultaneously.

30. The apparatus of claim 26, wherein when the new document matrix contains more new documents than new words, then the LSA space adapter transforms space by:

applying the word vector transformation matrix K, first; and

applying the document vector transformation matrix is J second, wherein the extension matrix Y is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Z.

31. The apparatus of claim 26, wherein when the new document matrix contains more new words than new documents, then the LSA space adapter transforms the LSA space by:

applying the document vector transformation matrix J first; and

applying the word vector transformation matrix K second, wherein the extension matrix Z is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Y.

32. The apparatus of claim 23, wherein the change in the language is a change in the language's domain.

33. The apparatus of claim 23, wherein the change in the language is a change in the language's style.

34. An apparatus for recognizing speech, the apparatus comprising:

means for recognizing an audio input as a new document; and

means for processing the new document using latent semantic adaptation; and

means, coupled to the means for processing, for semantically inferring from a vector representation of the new document which of a plurality of known words and known documents correlate to the new document.

35. The apparatus of claim 34, wherein the means for processing the sequence of words and documents using latent semantic adaptation comprises:

means for generating a latent semantic analysis (LSA) space from a training corpus of documents representative of a language;

means for receiving the new document that represents a change in the language; and

means for adapting the LSA space to reflect the change in the language.

36. The apparatus of claim 34, wherein the means for adapting the LSA space to reflect the change in the language comprises a means for transforming the LSA space to take into account the new document's influence on the LSA space without re-computing the LSA space.

37. The apparatus of claim 34, wherein the means for transforming the LSA space comprises:

means for obtaining a training document vector that characterizes a semantic position of the training document within the LSA space;

means for computing a new document vector that characterizes a semantic position of the new document within the LSA space;

means for deriving a document vector transformation matrix; and

means for applying the document vector transformation matrix to the training document vector and the new document vector to shift a position of each document vector in the LSA space, where the shift in the position reflects the change in the language.

38. The apparatus of claim 37, wherein the means for transforming the LSA space further comprises:

means for obtaining a training word vector that characterizes a semantic position of the training word within the LSA space;

means for computing a new word vector that characterizes a semantic position of the new word within the LSA space;

means for deriving a word vector transformation matrix; and

means for applying the word vector transformation matrix to the training word vector and the new word vector to shift a position of each word vector in the LSA space, where the shift in the position reflects the change in the language.

39. The apparatus of claim 38, wherein:

the training document vector is VS, where VS is computed from a right singular matrix V and a diagonal matrix S, each of which was obtained from a previous singular value decomposition (SVD) of a training word-document matrix constructed during the generation of the LSA space, the training word-document matrix representing the extent to which each of the words appears in each of the documents of the training corpus;

the new document vector ZS, where ZS is computed from the diagonal matrix S and an extension matrix Z, wherein Z is an extension of the right singular matrix V obtained by folding in a new word-document matrix, the new word-document matrix representing the extent to which a new word appears in the new document; and

the document vector transformation matrix is J, wherein J is obtained from a Choleski decomposition of a matrix derived from an extension matrix Y, wherein Y is an extension of a left singular matrix U obtained by folding in the new word-document

matrix, and wherein U was obtained from the previous SVD of the training word-document matrix constructed during the generation of the LSA space.

40. The apparatus of claim 39, wherein:

the training word vector is US, wherein US is computed from the left singular matrix U and the diagonal matrix S;

the new word vector is YS, where YS is computed from the the diagonal matrix S and the extension matrix Y; and

the word vector transformation matrix is K, wherein K is obtained from a Choleski decomposition of a matrix derived from the extension matrix Z.

41. The apparatus of claim 37, wherein the means for transforming the LSA space further comprises means for applying the document vector transformation matrix and the word vector transformation matrix simultaneously.

42. The apparatus of claim 37, wherein when the new document matrix contains more new documents than new words, then the means for transforming the LSA space further comprises:

means for applying the word vector transformation matrix K, first; and

means for applying the document vector transformation matrix J second, wherein the means for obtaining the extension matrix Y is not by folding in the new word-

document matrix, but is rather by deriving extension matrix Y from the extension matrix Z.

43. The apparatus of claim 37, wherein when the new document matrix contains more new words than new documents, then the means for transforming the LSA space further comprises:

means for applying the document vector transformation matrix J first; and

means for applying the word vector transformation matrix K second, wherein the means for obtaining the extension matrix Z is not by folding in the new word-document matrix, but is rather by deriving the extension matrix Z from the extension matrix Y.

44. The apparatus of claim 35, wherein the change in the language is a change in the language's domain.

45. The apparatus of claim 35, wherein the change in the language is a change in the language's style.

46. An system for processing speech, the system comprising:

a speech recognition database comprising a latent semantic analysis (LSA) space generated from a training corpus of documents representative of a language;

an input receiver to receive a new document that represents a change in the language; and

a processing system to adapt the LSA space to reflect the change in the language.

47. The system of claim 46, wherein the processing system adapts the LSA space by transforming the LSA space to take into account the new document's influence on the LSA space without re-computing the LSA space.

48. The system of claim 46, wherein the processing system transforms the LSA space by:

obtaining a training document vector that characterizes a semantic position of the training document within the LSA space;

computing a new document vector that characterizes a semantic position of the new document within the LSA space;

deriving a document vector transformation matrix; and

applying the document vector transformation matrix to the training document vector and the new document vector to shift a position of each document vector in the LSA space, where the shift in the position reflects the change in the language.

49. The system of claim 48, wherein the processing system further transforms the LSA space by:

obtaining a training word vector that characterizes a semantic position of the training word within the LSA space;

computing a new word vector that characterizes a semantic position of the new word within the LSA space;

deriving a word vector transformation matrix; and

applying the word vector transformation matrix to the training word vector and the new word vector to shift a position of each word vector in the LSA space, where the shift in the position reflects the change in the language.

50. The system of claim 49, wherein:

the training document vector is VS, where VS is computed from a right singular matrix V and a diagonal matrix S, each of which was obtained from a previous singular value decomposition (SVD) of a training word-document matrix constructed during the generation of the LSA space, the training word-document matrix representing the extent to which each of the words appears in each of the documents of the training corpus;

the new document vector ZS, where ZS is computed from the diagonal matrix S and an extension matrix Z, wherein Z is an extension of the right singular matrix V obtained by folding in a new word-document matrix, the new word-document matrix representing the extent to which a new word appears in the new document; and

the document vector transformation matrix is J, wherein J is obtained from a Choleski decomposition of a matrix derived from an extension matrix Y, wherein Y is an extension of a left singular matrix U obtained by folding in the new word-document matrix, and wherein U was obtained from the previous SVD of the training word-document matrix constructed during the generation of the LSA space.

51. The system of claim 50, wherein:

the training word vector is US, where US is computed from a left singular matrix U and the diagonal matrix S;

the new word vector is YS, wherein YS is computed from the diagonal matrix S and the extension matrix Y; and

the word vector transformation matrix is K, wherein K is obtained from a Choleski decomposition of a matrix derived from the extension matrix Z.

52. The system of claim 50, wherein the processing system transforms the LSA space by applying the document vector transformation matrix and the word vector transformation matrix simultaneously.

53. The system of claim 50, wherein when the new document matrix contains more new documents than new words, then the processing system transforms space by:

applying the word vector transformation matrix K, first; and

applying the document vector transformation matrix is J second, wherein the extension matrix Y is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Z.

54. The system of claim 50, wherein when the new document matrix contains more new words than new documents, then the processing system transforms the LSA space by:

applying the document vector transformation matrix J first; and

applying the word vector transformation matrix K second, wherein the extension matrix Z is not obtained by folding in the new word-document matrix, but is rather derived from the extension matrix Y.

55. The system of claim 46, wherein the change in the language is a change in the language's domain.

56. The system of claim 46, wherein the change in the language is a change in the language's style.